Knowledge-Driven Inference of Medical Interventions

Anthony Stell[1], Laura Moss[1,2], Ian Piper[1]
[1]*Department of Clinical Physics, University of Glasgow, Glasgow, UK*
[2]*Department of Computing Science, University of Aberdeen, Aberdeen, UK*
a.stell.1@research.gla.ac.uk

Abstract

Physiological monitoring equipment routinely collects large amounts of time series patient data. In addition to influencing the treatment of a patient, this data is often used in medical research. However, treatment data (e.g. sedation) can be difficult to collect. In this paper we describe the AMITIE (Automated Medical Intervention and Treatment Inference Engine) system which infers a medical intervention from physiological time series data. The system comprises several domain ontologies and an algorithm to detect abnormal physiological readings and infer the subsequent associated medical intervention. To evaluate this approach we have applied AMITIE in the neuro-intensive care unit domain.

1. Introduction

Time-series data is collected frequently across many fields of medicine. It is often the case that this data is used in medical research, for example to derive a surrogate measure of the efficacy of treatments for specific conditions. Medical interventions, defined as "*any examination, treatment, or other act having preventive, diagnostic, therapeutic, or rehabilitative aims and which is carried out by a physician or other health care provider*" [15] , are difficult to record, often because there is no general way to automate the capture of that information (c.f. passively monitoring vital signs by attaching equipment to the patient). In this paper we explore the inference of medical interventions from electronic data collected in the neuro-intensive care unit (ICU). In high-pressure environments such as these units, treatment is often administered and after the event noted in paper format, but not captured electronically. This lack of accessible treatment data is hindering research in this domain; in particular, in discovering which patient treatments are correlated with the best patient outcomes. To help solve some of these problems it is desirable to develop a system that can infer the presence of such interventions from medical data that is easily recordable, for example, from passively monitored vital signs of the patient. In this paper we introduce the AMITIE system. AMITIE uses detailed domain knowledge regarding an unknown intervention to infer its presence in the patient's data. To model this complex knowledge, several domain ontologies have been implemented. Inferences are then made using a set of general SPARQL queries [8].

The use of domain knowledge to generate abstractions (or inferences) from medical time series data is an established technique, enabling the intelligent analysis of data. For a review see [7]. Previous work describing the inference of medical interventions from time series data generally concentrates on identifying a particular trace in a data-set, e.g. identifying blood sampling from heart rate readings [10]. Some tools such as that developed by [1] allow a clinician to mark up features of interest in a data-set and, by applying a merging algorithm, identify patterns which represent an event. Other approaches have explored the modeling of clinical events associated with the target intervention. For example, in [2], suction episodes were inferred by identifying actions such as a change in ventilator settings. Temporal constraint networks were then used to represent the external action information (such as intervention administration or moving of patients) associated with the target. However, although these event detection methods perform well in each of their problem domains, the models are not easily transferred to other medical problems. For example, in [2], the approach requires the construction of a new temporal constraint network for every medical event. Additionally, relying on a physiological trace alone to identify an intervention can be difficult as some traces map to many interventions and medical conditions. The approach taken in the AMITIE system avoids some of these problems and allows for the potential reuse of AMITIE in other medical domains.

2. AMITIE

The AMITIE system (Automated Medical Intervention and Treatment Inference Engine) has been developed to identify abnormal physiological events and automatically infer subsequent medical interventions from time series physiological and treatment data. Figure 1 provides a high level overview of the system.

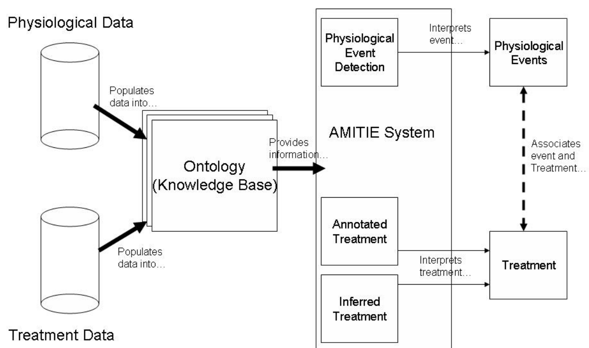

Figure 1: high-level overview of AMITIE system

The patient data is explored for instances of abnormal physiological readings. Once identified, the data is further examined to find related interventions given to the patient. In some cases the interventions are easy to identify as they have been specifically recorded in the patient data-set, with an explicit target noted. However, as described above interventions often have to be inferred. To enable an intervention to be inferred, detailed information can be obtained from domain ontologies. Information such as the known physiological effects of the intervention and other contextual information (e.g. contraindications of a drug) can help to determine whether it is likely that the patient has received the intervention for the abnormal reading. For example, consider a series of abnormally raised intracranial pressure readings which then return to normal. If the data-set is examined and no intervention is recorded, it may be reasonable to infer from the observation of the patient's temperature decreasing, that the procedure 'therapeutic cooling' has been administered.

In the AMITIE system, a set of ontologies are used to model the domain. An ontology "*defines a set of representational primitives with which to model a domain of knowledge or discourse*." [9] AMITIE's knowledge base consists of three OWL [13] domain ontologies which model the medical domain, patient data and human physiology. This knowledge base has been reused from previous work on the EIRA system [4]. The following high-level algorithm summarises the functionality of the AMITIE system:

1. Identify an abnormal physiological event (E)
   a. Characterize E into whether it has returned to a baseline ("normal") value or not.
   b. If it has, it is assumed that the patient has been treated and the time period (TP) of E is determined.
   c. Examine TP for instances of annotated interventions
      i. If intervention (I) is noted in data-set: suggest that I has been given in response to E.
      ii. Else, infer non-annotated intervention:
      iii. Identify known physiological effects of possible interventions for E.
      iv. Examine TP for evidence of any of these effects.
      v. If possible intervention (I_p) is found:
      vi. suggest I_p has been given in response to E
      vii. Else, suggest that patient returned to baseline value without intervention.

Issues of negation and ranking have not yet been handled. This algorithm also assumes a one-to-one relationship between event and physiological output. However, it is likely that combinations and emergent factors that make this relationship complex.

3. Implementation

The AMITIE system software is written using the JENA API [11] (in Java 6 [12]). To obtain the information from the ontology, a set of SPARQL queries have been implemented. These allow for a separation of the inference process from the application code (for later re-use in other medical contexts). They have the following functions and features:

1) The first query obtains all the physiological data above a certain threshold:

```
SELECT ?timepoint ?physiovalue
WHERE {?x <http://www.owl-
ontologies.com/amitie.owl#hasTime>?timepoint .
[...]
?reading <http://www.owl-
ontologies.com/amitie.owl#readingParameter>
<http://www.owl-
ontologies.com/amitie.owl#"[physiolabel]"> .
FILTER (?physiovalue > "[physioThreshold]"}.}
```

Of the variables in the query, "physiolabel" refers to the physiological parameter that is being queried (e.g. heart rate) and "physioThreshold" refers to the value above which an abnormal event is deemed to have occurred (e.g. ICP readings above 20 mmHg are generally considered to require treatment).

2) The second query obtains all the treatments that *are* annotated in the data-set within the time period of a single event.

3) The information required to infer interventions, for when a treatment is not annotated in the data-set, is extracted from the ontology using five queries that interrogate the following features to see if they are present (as they are significant in the event signature):

- *High_Feature*,
- *Low_Feature*,
- *Increase_Parameter_Change*, *Decrease_Parameter_Change*
- *Constant_Parameter*.

4) The final query obtains other physiological data that may relate to the abnormal event in question, but has not been retrieved in the original query. The layout is similar to the first query, but is constrained by time series rather than threshold physiological values.

The other main part of the AMITIE system is the detection of interventions using the data that has been retrieved from the domain ontology. The implemented procedure followed the algorithm outlined in section 2.2. An example of this would be to search for increased systolic blood pressure and constant heart rate, which would indicate that vasopressors have been administered. This procedure is implemented in a TreatmentCheck class and is passed back to the main program for each event and each listed potential intervention. Using this combination of domain knowledge and patient data, the resulting management of a patient's abnormal event is determined to either be an annotated treatment, a non-annotated treatment, or the patient's vital sign has returned to normal without any clinical intervention.

4. Evaluation: detecting interventions for the management of raised intracranial pressure (ICP)

To evaluate the AMITIE system, it has been applied to data from the neuro-intensive care domain to identify interventions following a period of raised intracranial pressure (ICP) readings. Raised ICP is a frequent consequence of traumatic brain injury where injury related mechanisms such as intracranial bleeding or ischaemia can cause brain tissue swelling which, in an enclosed container like the skull, will produce rapid life threatening increases in intracranial pressure. Raised ICP is typically defined as an ICP level raised above 20 mmHg, for a sustained period of several minutes [14]. It is useful for researchers to know what treatment was given following a raised ICP so that it can be compared with patient outcomes.

4.1 Data-sets: Brain-IT and Avert-IT

The patient data that the AMITIE system has been evaluated on has been collected as part of the Brain-IT [5] and Avert-IT [6] projects, as described by the core Brain-IT data standard [3]. The data-set has a comprehensive coverage (72%) and is one of only a few data-sets that include linked clinical outcome data (Glasgow Outcome Scale at 6 months post-injury) to the treatments and vital signs monitored during the patient's stay in the ICU [3].

4.2 Intervention Signatures

With input from clinical domain experts, the physiological signatures of four important neurological ICU interventions, likely to be applied for raised ICP events, were obtained (table 1).

Treatment	Signature
Ventilation	Altered Oxygenation (SaO2) OR respiration rate (RR) OR blood gas (PaCO2)
Induced hypothermia	Lowered Temperature (Tc)
Vasopressors (cerebral vasoconstriction)	Increased mean arterial pressure (MAP) AND constant heart rate (HR)
Sedative reduction	Increased mean arterial pressure (MAP) AND increased heart rate (HR)

Table 1: Treatments for ICP management and their physiological signature

4.3 Method

The following steps were applied in this case study:

- Initially Brain-IT data was converted into RDF. For the study we randomly selected three patients consisting of 1000, 5000, 13,000 time points and 25, 209, and 77 raised ICP events respectively.
- The physiological signatures of the possible interventions were obtained from domain experts and used to add additional instances, specific to a neuro-ICU, to the ICU domain ontology (e.g. cerebrospinal fluid drainage and cerebral vasoconstriction).

- The AMITIE system was then run, and instances of annotated, inferred, or absent treatments were obtained for each noted ICP event.

4.4 Results

Table 2 shows the results of running the AMITIE system against a total of 311 raised ICP events.

Patient ID	No. of ICP events	Annotated Treatments for each ICP event	Inferred Treatments for each ICP event
15026161	25	0	Ventilation and cerebral vasoconstriction
15127262	209	0	Ventilation and cerebral vasoconstriction
15137626	77	0	Ventilation, induced hypothermia and cerebral vasoconstriction

Table 2: AMITIE results

AMITIE successfully produced an inferred treatment for all of the 311 ICP events. Within the context of an individual patient, the same treatments were suggested for all of the ICP events. This is likely to be a result of clinicians following protocol guidelines, or at least interventions consistent for a single patient so perhaps should not be unexpected. For each patient the annotated interventions associated with the ICP events recorded are zero. Again as discussed earlier, this should not be unexpected as interventions are often not recorded accurately. Further investigation of the data showed that some interventions were annotated near these events (but out with the analysed time period); a subset of these matched the inferred interventions suggested by AMITIE, this is very promising.

5. Discussion & Future Work

This paper outlines the early stages of an approach to inferring clinical interventions events from time series data using domain ontologies. The approach described in this paper exploits the interaction of relationships as modeled in the supporting knowledge base. This is distinct from, and improves upon, other work in the area by using qualitative relationship information, rather than calculated pattern relationships. The evaluation has produced promising results and a further evaluation of the tool on a greater number of patients is planned. Care must be taken when making medical interpretations of physiological output (e.g. low temperature levels indicating a patient's hypothermic state rather than artificial application as an intervention). More contextual information is required in the knowledge base to discriminate between such situations. Planned future work includes an evaluation on a wider variety of patients available from the Brain-IT data-set (containing more than 200 patients and a further 60 from the Avert-IT project). A wider variety of patients will expose AMITIE to different types of treatments that will need to be inferred. In addition, a study is planned in which we will attempt to validate the accuracy of the intervention inference. Although the inferences made in the initial evaluation are clinically sensible, they have not been validated against what *actually* happened to the patient. To do this we plan to compare the inferred intervention against a known annotated intervention (either recorded electronically or in the patient notes). This would also allow the independent verification of the intervention signatures outlined in section 4.2.

Acknowledgements The authors would like to acknowledge the work of the BrainIT group investigators and participating centres to the BrainIT data-set without whom this work could not have been conducted [5].

6. References

[1] Hunter, McIntosh, et al – Knowledge-based event detection in complex time-series data, *Artificial Intelligence in Medicine 1999*, p271-280
[2] Feng G, et al – Using temporal constraints to integrate signal analysis and domain knowledge in medical event detection, *Artificial Intelligence in Medicine 2009,* p46-55
[3] Enblad P et al - *R3 Survey of traumatic brain injury management in European Brain IT centres year 2004*, Intensive Care Med (2004) 30:1058-1065
[4] Moss L - *Explaining Anomalies: An Approach to Anomaly-Driven Revision of a Theory*, PhD thesis, University of Aberdeen 2010
[5] Brain-IT – http://www.brainit.eu
[6] Avert-IT – http://www.avert-it.org
[7] Stacey, McGregor, et al – Temporal abstraction in intelligent clinical data analysis: a survey. *Artificial Intelligence in Medicine 2007*, p1-24
[8] SPARQL - http://www.w3.org/TR/rdf-sparql-query/
[9] Tom Gruber, *Ontology - Entry in the* Encyclopedia of Database Systems, Springer-Verlag, 2008.
[10] Williams, Stanculescu - Automating the calibration of a neonatal condition monitoring system. 13th Conference, *Artificial Intelligence in Medicine 2011*, p240-249.
[11] JENA - http://incubator.apache.org/jena/
[12] Java 6 – http://www.java.com
[13] OWL - http://www.w3.org/TR/owl-ref/
[14] Steiner LA, Andrews PJ (2006). *Monitoring the injured brain: ICP and CBF*. British Journal of Anaesthesia 97 (1): 26–38.
[15] WHO - *Declaration on the Promotion of Rights of Patients in Europe*, Amsterdam 1994